Simulation Summarization

(Extended Abstract)

Nidhi Parikh, Madhav V. Marathe, Samarth Swarup
Network Dynamics and Simulation Science Lab,
Biocomplexity Institute of Virginia
Virginia Tech,
Blacksburg, VA, USA.
{nidhip,mmarathe,swarup}@vbi.vt.edu

ABSTRACT

As increasingly large-scale multiagent simulations are being implemented, new methods are becoming necessary for concisely summarizing the results of a simulation run. Here we pose this as the problem of simulation summarization: how to extract the causally-relevant states from the trajectories of the agents. We present a simple algorithm to compress agent trajectories through state space by identifying the state transitions which have significant impact on the final outcome of interest. We apply it to a complex simulation of a major disaster in an urban area and present results.

Keywords

simulation summarization, causal states

1. INTRODUCTION

Large-scale multiagent simulations are becoming increasingly common in many domains of scientific interest, including economics, epidemiology, social science, and disaster response. These simulations have complex models of agents, environments, infrastructures, and interactions.

Often the goal is to study a hypothetical situation in a detailed and realistic virtual setting, with the intention of making policy recommendations. However, analyzing results for such simulations pose some challenges:

1. They are computationally too expensive to run enough number of times to obtain the statistical power.
2. Often the kind of interventions to evaluate are not known. Ideally, we would like to run simulation to get some insights that help suggest interventions.
3. Large-scale simulations can generate much more data in each simulation run than goes into the simulation.

As a first step towards addressing these kinds of problems, we introduce the problem of simulation summarization. Our perspective on summarizing a multiagent simulation is that the summary should capture the causally-relevant states of the simulation. By "causally-relevant", we mean agent states that have a measurable impact on outcomes of interest, even if the impact is delayed. In line with this intuition, we adapt the approach of "causal states" for stochastic process.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

2. CAUSAL STATES

Consider a stochastic process as a sequence of random variables X_t, drawn from a discrete alphabet, $\mathcal{A}$. Let $\overleftarrow{X} = X_{-\infty} \dots X_{t-2}X_{t-1}X_t$ and $\overrightarrow{X} = X_{t+1}X_{t+2} \dots X_{\infty}$ denote the *past* and the *future* of the sequence, respectively.

Crutchfield et. al. [1] suggested a simple method for modeling a stochastic process that captures the mutual information that is communicated from the past $\overleftarrow{X}$ to the future $\overrightarrow{X}$ of the sequence: group together all the histories that predict the same future. This gives rise to a state machine, called an ϵ-machine, defined as:

$$\epsilon(\overleftarrow{x}) = \{\overleftarrow{x}' | Pr(\overrightarrow{X}|\overleftarrow{x}) = Pr(\overrightarrow{X}|\overleftarrow{x}')\}. \quad (1)$$

Shalizi et. al. have presented an algorithm for learning ϵ-machine from time series, known as *Causal State Splitting Reconstruction* (CSSR) [4]. It learns an ϵ-machine as a Hidden Markov Model (HMM) in an incremental fashion. The HMM is initialized with one state which consists of null history of length 0. The algorithm tests the distribution over the next symbol given progressively longer past sequences. Let L be the length of the past sequences considered so far. At the next step, it looks at sequences of length $L+1$, e.g., ax^L where x^L is a sequence of length L and $a \in \mathcal{A}$. If ax^L belongs to the same causal state as x^L, then we would have,

$$Pr(X_t|ax^L) = Pr(X_t|\hat{S} = \hat{\epsilon}(x^L)), \quad (2)$$

where $\hat{S}$ is the current estimate of the causal state to which x^L belongs. This hypothesis can be tested using a statistical test such as the Kolmogorov-Smirnov test. If the test shows that the LHS and RHS of equation 2 are statistically significantly different distributions, then CSSR tries to match the sequence ax^L with all the other causal states estimated so far. If $Pr(X_t|ax^L)$ turns out to be significantly different in all cases, a new causal state is created and ax^L is assigned to it. This process is carried out up to some length L_{max}.

3. OUR APPROACH

Our approach adapts the causal state formalism by treating the trajectory of each agent in the simulation as an instance of the same stochastic process. Let N be the number of agents in the simulation. State of an agent a is denoted by a k-dimensional state vector $\mathbf{x_a}(t) = [x_1(t), x_2(t), \dots x_k(t)]^\intercal$, which evolves over time. Let d_i be the number of possible values x_i can take and hence $d = d_1 \times d_2 \times \dots \times d_k$ is the number of ways in which agent state can evolve. Let the outcome variable for agent a be denoted by y_a.

Our goal is to compress the trajectory of each agent through state space to a small number of important states that have a significant impact on the outcomes we care about.

At each time step t, the agent population is devided into a set of clusters, $C(t) = \{C_1(t) \cup C_2(t) \cup \ldots C_m(t)\}$. Initially, all the agents belong to one cluster. At each next time step, the state of each agent can change in d ways. Hence, an arbitrary cluster of agents, $C_i(t)$ can split into up to d groups at time step $t+1$. But not all of these changes may have a significant impact on the outcome. We treat each group derived from $C_i(t)$ as a candidate cluster, $CC_{i,j}(t+1)$, where $j \in 1 \ldots d$. At each step, we compare $Pr(Y|C_i(t))$ with $Pr(Y|CC_{i,j}(t+1)$ using the Kolmogorov-Smirnov test. Our null hypothesis (analogous to equation 2) is,

$$Pr(Y|CC_{i,j}(t+1)) = Pr(Y|C_i(t)). \quad (3)$$

We also introduce a parameter δ, which is a threshold on the "effect size", measured as the KL-divergence between $Pr(Y|C_i(t))$ and $Pr(Y|CC_{i,j}(t+1))$. If the null hypothesis is rejected at a level α (say 0.001) and $D_{KL}(Pr(Y|C_i(t))|| Pr(Y|CC_{i,j}(t+1))) > \delta$, then candidate cluster $CC_{i,j}(t+1)$ is accepted as a new cluster at time $t+1$. If none of the candidate clusters at time step $t+1$ are accepted, then $C_i(t)$ is added to the set of clusters for time step $t+1$.

Thus, the entire simulation is decomposed into a tree structure of agent clusters. Furthermore, each cluster splits only when the corresponding state change is informative about the final outcome of concern. The trajectory of each agent traces a path through this tree structure. We compress the trajectory by retaining only those time steps at which the cluster to which the agent belongs splits off from its parent cluster. The set of compressed agent trajectories constitutes our summary representation of the simulation.

The parameter δ allows to control the number of new clusters formed and consequently, the amount of compression. A high value of δ will retain only the clusters which have a large difference in outcomes from their parent clusters.

4. EXPERIMENTS

We apply our algorithm to a multiagent simulation of a hypothetical improvised nuclear device detonation in Washington DC [3]. Our simulation consists of a large, detailed "synthetic information system" [2] which represents the human population of the region and detailed models of four infrastructures: cell phone communication system, power system, transportation system, and healthcare system.

Agents are defined by a number of state variables. However, for the purpose of summarization, we focus on six variables: health (modeled on a 0 to 7 range where 0 is dead and 7 corresponds to full health), behavior (six behaviors mentioned below, plus categories indicating if agent is in healthcare location or out of the affected area), whether the agent has received an emergency broadcast (EBR), the agent's exposure to radiation, whether the agent has received treatment, and the agent's distance from the ground zero.

Agent behavior is conceptually based on the formalism of decentralized semi-Markov decision process (Dec-SMDP) with communication using the framework of options. High level behaviors are modeled as collection of options. We model six behaviors: household reconstitution (HRO), evacuation, shelter-seeking, healthcare-seeking, panic, and aid & assist. These high level behavior options are policies over low level actions which are to call, text or move. These actions are supported by infrastructural systems. Details of agent design and behavior can be found in [3].

4.1 Results

Due to limitation of space, we only present results from one value of KL-divergence threshold, $\delta = 5$. Compressed trajectories of agents are stored in a database table along with the expected value of final health state which can be used to query any subpopulation of interest.

We evaluate two queries that identify top 10 transitions in first three hours where current health state remains same (so improvement is not due to current health state) but the expected health state is improved or reduced, ordered by expected improvement or reduction in descending order.

Results show that for agents who are close to ground zero and in health state 4 or 5, reaching healthcare location early helps to improve health, even if the exposure level is high. For agents in healthstate 5 to 7, who are far from ground zero though with medium exposure, getting out of the area helps. Also, for agents in health state 7, who are far from ground zero, with medium exposure, and panicking, receiving EBR helps (as it provides information about the event and recommends sheltering).

For agents who are currently in a full health, close to ground zero, and have high exposure level, doing household reconstitution (HRO) reduces their expected outcome. Even if the current health state is good, this accounts for the delayed effect of radiation. For people who are already in low health (health state 3), panicking or seeking healthcare (which makes them travel to healthcare location and get exposed to more radiation) deteriorates expected health state, even if far from ground zero.

Acknowledgments

This work has been supported in part by DTRA CNIMS Contract HDTRA1-11-D-0016-0001, DTRA Grant HDTRA1-11-1-0016, NIH MIDAS Grant 5U01GM070694-11, NIH Grant 1R01GM109718, NSF NetSE Grant CNS-1011769, and NSF SDCI Grant OCI-1032677.

REFERENCES

[1] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys Rev Lett*, 63(2):105–108, 1989.

[2] M. Marathe, H. Mortveit, N. Parikh, and S. Swarup. Prescriptive analytics using synthetic information. In W. H. Hsu, editor, *Emerging Trends in Predictive Analytics: Risk Management and Decision Making.* IGI Global, 2014.

[3] N. Parikh, S. Swarup, P. E. Stretz, C. M. Rivers, B. L. Lewis, M. V. Marathe, S. G. Eubank, C. L. Barrett, K. Lum, and Y. Chungbaek. Modeling human behavior in the aftermath of a hypothetical improvised nuclear detonation. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Saint Paul, MN, USA, May 2013.

[4] C. R. Shalizi and K. L. Shalizi. Blind construction of optimal nonlinear recursive predictors for discrete sequences. In M. Chickering and J. Halpern, editors, *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence*, pages 504–511, Banff, Canada, 2004.